

00/03/00



JC760 U.S. PTO

Date: March 3, 2000File No. A-66977/JAS/RMA/LM

**Box PATENT APPLICATION FEE**  
 Assistant Commissioner for Patents  
 Washington, DC 20231-0001

"EXPRESS MAIL" MAILING LABEL NUMBER EL406568229USDATE OF DEPOSIT March 3, 2000

I HEREBY CERTIFY THAT THIS PAPER OR FEE IS BEING DEPOSITED WITH THE UNITED STATES POSTAL SERVICE "EXPRESS MAIL POST OFFICE TO ADDRESSEE" SERVICE UNDER 37 CFR 1.10 ON THE DATE INDICATED ABOVE AND IS ADDRESSED TO: BOX PATENT APPLICATION FEE, ASSISTANT COMMISSIONER FOR PATENTS, WASHINGTON, DC 20231-0001.  
 TYPED NAME Grace de Bos

SIGNED

Sir:

Transmitted herewith for filing is the patent application of Inventor(s): **Govindaraju GNANASIVAM** and **Krishnakumar Rao SURUGUCCHI**

For: **APPARATUS AND METHOD FOR IMPLEMENTING RAID DEVICES IN A CLUSTER COMPUTER SYSTEM**

Enclosed are also:

Prior Art Statement  
 29 Sheets of drawing, Formal       , Informal   
 Recordation Form Cover Sheet  
 An Assignment of the invention to: Mylex Corporation  
 Declaration for Patent Application  
 Power of Attorney  
 Small Entity Status Declaration Under 37 CFR 1.9(f) and 1.27(c)

JC564 U.S. PRO  
 09/518551  
 03/03/00

	(Col. 1) NO. FILED	(Col. 2) NO. EXTRA	SMALL ENTITY RATE	OTHER THAN SMALL ENTITY RATE
<b>BASIC FEE</b>			\$345	\$690
TOTAL CLAIMS	20 - 20 =	0	x 9 = \$	x 18 = \$
INDEP CLAIMS	6 - 3 =	3	x 39 = \$	x 78 = \$234
MULTIPLE DEPENDENT CLAIM PRESENTED			+130 = \$	+260 = \$
If the difference in Col 1 is less than zero, enter "0" in Col. 2			TOTAL \$	TOTAL \$924

Our Check No. 28719 in the amount of \$924.00 is enclosed to cover the filing fee.

The Commissioner is hereby authorized to charge any additional fees which may be required, including extension fees, or credit any overpayment to Deposit Account No. 06-1300 (Order No. A-66977/JAS/RMA/LM).

Two copies of this sheet are enclosed.

Respectfully submitted,

Larry Mendenhall  
 Reg. No. 38,555

FLEHR HOHBACH TEST ALBRITTON & HERBERT LLP  
 Four Embarcadero Center, Suite 3400  
 San Francisco, California 94111-4187  
 Telephone: (415) 781-1989  
 Fax: (415) 398-3249  
 1011617

5

PATENT APPLICATION

10

APPARATUS AND METHOD FOR IMPLEMENTING RAID DEVICES IN A  
CLUSTER COMPUTER SYSTEM

15

Inventors:

20

Govindaraju Gnanasivam  
a citizen of India, residing at  
1189 Moulton Drive  
Milpitas, California 95035

25

Krishnakumar Rao Surugucchi  
a citizen of India, residing at  
810 San Marco Avenue  
Fremont, California 94539

30

35

Assignee:

MYLEX CORPORATION,  
a Delaware Corporation  
34551 Ardenwood Blvd.  
Fremont, California 94555-3607

40

FLEHR HOHBACH TEST ALBRITTON & HERBERT LLP  
4 Embarcadero Center  
Suite 3400  
San Francisco, CA 94111-4187

45 (415) 781-1989

5

# APPARATUS AND METHOD FOR IMPLEMENTING RAID DEVICES IN A CLUSTER COMPUTER SYSTEM

## CROSS-REFERENCE TO RELATED APPLICATIONS

This application relates to:

10 U.S. Patent Application No. 09/\_\_\_\_, entitled, "Apparatus and Method for Rebuilding a Logical Device in a Cluster Computer System," filed \_\_ \_\_, 2000, naming Govindaraju Gnanasivam and Nallakumar Meiyappan as inventors, with Attorney Docket No. A-66978/JAS/RMA/LM / No. 38, and commonly assigned to Mylex Corp. of Fremont, CA;

15 U.S. Patent Application No. 09/\_\_\_\_, entitled, "Apparatus and Method for Detecting the Reset of a Node in a Cluster Computer System," filed \_\_ \_\_, 2000, naming Mohan B. Rowlands as inventor, with Attorney Docket No. A-66997/JAS/RMA/LM / No. 43, and commonly assigned to Mylex Corp. of Fremont, CA; and

20 U.S. Patent Application No. 09/\_\_\_\_, entitled, "Suicide Among Well-Mannered Cluster Nodes Experiencing Heartbeat Failure," filed \_\_ \_\_, 2000, naming Mohan B. Rowlands as inventor, with Attorney Docket No. A-67594/JAS/RMA/LM / No. 65, and commonly assigned to Mylex Corp. of Fremont, CA.

25

## BACKGROUND

This invention relates to cluster computer systems in general. More particularly, the invention relates to recovering from cable failure in cluster computer systems with RAID devices.

30

Historically, data-center operators running critical applications demanding high reliability have turned to mainframes, minicomputers and the like running complex fault-tolerant software on complex fault-tolerant hardware. In a different market niche of less critical and less demanding environments, the Microsoft Corp. Windows operating system has made significant inroads into business data centers, running on relatively inexpensive and uncomplicated personal-computer and server platforms. These Windows platforms were adequate for certain services – database and e-mail, for example.

10 However, databases and e-mail are becoming increasingly important in the average business. Indeed, in some businesses these functions have taken on a critical nature. Accordingly, data-center operators with now-critical database, e-mail and similar applications want to run them on systems with high reliability. They are unwilling, 15 however, to pay the costs of mainframes, minicomputers and their fault-tolerant software. In response to market demand, Microsoft Corp. has modified its Windows operating system to address the issue of high reliability.

20 Specifically, Microsoft now offers a Cluster Service product. Venerable if not ancient in the art, a “cluster” can be loosely defined as a parallel or distributed system of interconnected whole computers (sometimes called “systems” but herein termed “nodes” for clarity). The user of a cluster system logically views and uses it as a single, unified computing resource or service.

25 Generally speaking, a cluster enables the sharing of a computing load over several nodes without the user or client needing to know that more than one constituent node is involved. If any hardware or software component in the cluster system fails, the user or client may notice degraded performance but does not lose access to the service. The 30 cluster system disperses the load from the failed component to the remainder of the cluster system. Conversely, if the user or client notices the need for more of a given resource (for example, processing

power), that resource is simply added to the running cluster system, and the performance of the cluster system as a whole improves.

Well known in the art and only generally described here, the Microsoft Cluster Service product is the collection of all cluster-  
5 activity-management software on each node of a Microsoft cluster system. The Cluster Service is more fully described in "Microsoft Windows NT Server Cluster Strategy: High Availability and Scalability with Industry-Standard Hardware," (Microsoft Corp., 1995) and "Concepts and Planning: Microsoft 'Wolfpack' Clustering for Windows NT Server" (Microsoft Corp.,  
10 19\_\_). These two Microsoft clustering documents are attached hereto as Appendices A and B and are incorporated by reference as well.

A Microsoft cluster system uses the Small Computer Systems Interface (SCSI) bus with multiple initiators as the storage connection (although Microsoft envisions supporting the Fiber Channel in the  
15 future). Well known in the art, SCSI is an evolving standard directed toward the support of logical addressing of data blocks on data devices. Documents detailing the variations of SCSI over time (SCSI-1, SCSI-2 and SCSI-3, for example) are available from the American National Standards Institute (ANSI) of New York, New York ([www.ansi.org](http://www.ansi.org)). SCSI-1, SCSI-2  
20 and SCSI-3 are together referred to as "SCSI" herein.

Figure 1 illustrates a two-node cluster system 100 implemented on a SCSI bus 110 according to the prior art. In Figure 1, the cluster system 100 includes a first server node 120a and a second server node 120b. The server nodes 120a and 120b have respective SCSI  
25 identifiers (SCSI IDs) 7 and 6. The server nodes 120 connect to the SCSI bus 110 through respective host bus adapters (HBAs) 121.

A node 120 typically includes one or more of the following: a central processor unit ("CPU") 126, a memory 122, a user interface 123, a co-processor 124, ports 125, a communications interface 121 and an  
30 internal bus 127.

Of course, in an embedded system, some of these components may be missing, as is well understood in the art of embedded systems. In

a distributed computing environment, some of these components may be on separate physical machines, as is well understood in the art of distributed computing.

The memory 122 typically includes high-speed, volatile 5 random-access memory (RAM) 1221, as well as non-volatile memory such as read-only memory (ROM) 1223. Further, the memory 122 typically contains software 1222. The software 1222 is layered: Application software 12221 communicates with the operating system 12222, and the operating system 12222 communicates with the I/O subsystem 12223. The I/O subsystem 12223 10 communicates with the user interface 123, the co-processor 124 and the communications interface 121 by means of the communications bus 127.

The communications interface 121, in this embodiment, is a host bus adapter 121.

The communications bus 127 communicatively interconnects the 15 CPU 126, memory 122, user interface 123, co-processor 124 and communications interface 121.

To the SCSI bus 110 are also connected SCSI devices 130. The devices 130a through 130c can be, for example, physical disks with SCSI 20 IDs 0 through 2, respectively.

Local disks 150 connect to respective nodes 120 as necessary.

Figure 20 illustrates the physical view of a second cluster system 2000 implemented on a SCSI bus 110 with an external RAID controller 2060, according to the prior art. As in the cluster system 100, the cluster system 2000 includes the first and second server nodes 25 120. The server nodes 120 have respective SCSI IDs 7 and 6 and connect to the SCSI bus 110 through respective HBAs 121. Each of the nodes 120 runs software 1222.

To the SCSI bus 110 is also connected the device 130a and a RAID controller 2060 with respective unique SCSI IDs. Additional SCSI 30 devices 2061 attach to the RAID controller 2060 by means of a SCSI bus 2062. The devices 130, 2061 can be physical disks, for example.

Again, local disks 150 connect to respective nodes 120 as

necessary.

Figure 21 illustrates the logical view of the cluster system 2000 of Figure 20. The device 130 and the RAID controller 2060 each appears to the host 120 as a single SCSI device. The RAID controller 5 2060 organizes the devices 2061 to appear to the host 120 as logical units (LUNs) 2063 of the SCSI device 2060.

Figure 22 illustrates the physical view of a third cluster system 2200 with internal RAID controllers 2210 and multiple shared SCSI channels 110, according to the prior art. As in the previous systems, 10 the cluster system 2200 includes the first and second server nodes 120 with respective SCSI IDs 7 and 6. The server nodes 120 connect to multiple SCSI buses 110 through respective RAID controllers 2210 and run the software 1222.

To each SCSI bus 110 is connected at least one device 2061, 15 each device 2061 having a SCSI ID unique for the channel 110 to which it connects. The devices 2061 can be physical disks, for example. Local disks 150 again connect to respective nodes 120 as necessary.

Figure 23 illustrates the logical view of the cluster system 2200 of Figure 22. The RAID controllers 2210 organize the devices 2061 20 to appear to the host 120 as SCSI disks 130 on a single SCSI channel 110. The RAID controllers 2210 thus appear to the host 120 as HBAs 121.

In this sense, the RAID controllers 2060, 2210 hide the complexity of the RAID SCSI disks 2061 and the controllers 2060, 2210 themselves from the hosts 120.

25 The SCSI standard implements Reserve() and Release() commands. This pair of commands allows a SCSI initiator (for example, a node 120) to reserve a SCSI target or logical unit on a SCSI target and later to release it. In the prior art, the usual handling of one of these commands in a Microsoft Windows 95/98 cluster system 100, 2000, 30 2200 involves an HBA 121 passing the command to the target, which then executes it.

Where the SCSI target of a Reserve() command is a logical

unit 2063 of an external RAID controller 2060 or where the SCSI target is a logical disk 130 depending from an internal RAID controller 2210, the controller 2060, 2210 still passes the Reserve() command to all of the disks 2061 that compose the target. This pass-through method, however, 5 is patently inefficient, reserving more devices 2061 than the initiator 120 may require. The pass-through method also imposes limitations on a RAID configuration.

The implementation of a RAID device 2060 in a cluster environment presents another problem, this with respect to disk failure. 10 In a non-cluster environment, rebuilding a logical device 2063, 130 in the face of failure is a well-practiced art: A controller restores data from a mirroring physical drive to a replacement physical drive. In a non-cluster environment, the logical choice of which node 120 is to rebuild the failed logical device 2063, 130 is the one and only node 120 15 holding the reservation to any of the physical units 2061.

In a cluster environment, however, multiple nodes 120 can hold a reservation to a physical unit 2061 through reservations to logical devices 2063, 130 comprising that unit 2061. Further, one node 120 can reserve a logical device 2063, 130 while a different node 120 20 receives the command to rebuild the logical device 2063, 130.

Accordingly, it is desirable to handle more efficiently and less restrictively the SCSI Reserve() and Release() commands in a cluster environment with RAID devices.

Also, in a cluster environment with RAID devices, it is 25 desirable to rebuild a logical unit in a manner simple and localized to the affected nodes.

These and other goals of the invention will be readily apparent to one of skill in the art on reading the background above and the description below.

## SUMMARY

Herein are described apparatus and methods for efficiently implementing logical-device reservations in a cluster computer system.

5 The apparatus include cooperating controllers programmed in firmware around a distributed reservation table.

In one embodiment, an apparatus manages access to an I/O device, with first and second nodes with respective bus controllers communicatively coupled to each other and to a logical I/O device by 10 means of a bus. The first controller receives a request to reserve the logical device and, in response, communicates a reservation request for the logical device over the bus to the second controller for execution by the second controller.

In another embodiment, in response to the communicated 15 reservation request the second controller reserves the logical device for the first node.

In another embodiment, the logical device is a logical device depending from a multi-logical-device controller.

In another embodiment, before communicating the reservation 20 request, the first controller determines whether the logical device is already reserved within the first controller and aborts managing the access when the logical device is already reserved. Otherwise, the first controller reserves the logical device for the first node in response to the reservation request.

25 In another embodiment, the invention further manages access by the first controller's receiving a request to release the logical device and responsively communicating a release request over the bus to the second controller for execution by the second controller. The second controller responsively releases the logical device within itself.

## BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a two-node cluster system implemented on a SCSI bus according to the prior art.

5 Figure 2 illustrates the physical view of a two-node cluster computer system according to one embodiment of the invention.

Figure 3 illustrates a controller for a node of a cluster system according to the invention.

10 Figure 4 is a flowchart illustrating the startup protocol for a master node.

Figures 5A-5C compose a flowchart illustrating the startup protocol for a slave node.

Figure 6 is a state diagram giving the states transitions a master node makes regarding a slave node.

15 Figure 7 is a state diagram giving the state transitions a slave node makes regarding itself.

Figures 8A-8C compose a flowchart illustrating actions of the master node's monitor process.

20 Figure 9 illustrates the logical view of the cluster system 200 of Figure 2.

Figure 10 is a flowchart illustrating flow of control in a node processing a SCSI Reserve() command.

Figures 11A-11C compose a flowchart illustrating flow of control in a node processing SCSI bus communications.

25 Figure 12 is a flowchart illustrating flow of control in a node processing a Node\_Reserve() command.

Figure 13 is a flowchart illustrating flow of control in a node processing a Node\_Negotiate() command.

30 Figure 14 is a flowchart illustrating flow of control in a node controller receiving a rebuild command.

Figure 15 is a flowchart illustrating flow of control in a node controller confirming a reported configuration change.

Figure 16 is a flowchart illustrating flow of control in a master node controller on receipt of a Node\_Inform\_Master\_Rebuild\_Received() command.

5 Figure 17 is a flowchart illustrating flow of control in a node controller monitoring another controller.

Figure 18 is a flowchart illustrating flow of control in a node controller processing a Node\_Rebuild() command.

Figure 19 is a flowchart illustrating flow of control in a node controller processing a SCSI\_Release() command.

10 Figure 20 illustrates the physical view of a second cluster system implemented on a SCSI bus with an external RAID controller, according to the prior art.

Figure 21 illustrates the logical view of the cluster system of Figure 20.

15 Figure 22 illustrates the physical view of a third cluster system with internal RAID controllers and multiple shared SCSI channels, according to the prior art.

Figure 23 illustrates the logical view of the cluster system of Figure 22.

20

## DESCRIPTION OF SPECIFIC EMBODIMENTS

### TABLE OF CONTENTS

Overview . . . . .	11
Data Structures . . . . .	13
25 Methods, Processes and Protocols . . . . .	21
Firmware-Implemented SCSI Commands . . . . .	21
Node-to-Node Communications . . . . .	22
Node Startup . . . . .	27
The Monitor Process . . . . .	30
30 Configuration Changes . . . . .	35
Rebuild of a Logical Device . . . . .	36
WHAT IS CLAIMED IS . . . . .	58

## Overview

Figure 2 illustrates the physical view of a two-node cluster computer system 200 according to one embodiment of the invention. In 5 Figure 2, the computer system 200 includes nodes 220a and 220b. The server nodes 220 have respective SCSI identifiers (SCSI IDs) 7 and 6 (on all SCSI buses 110, though this need not be true). The server nodes 220 connect to SCSI buses 110 through SCSI controllers 221. SCSI controllers 221 connected to the same SCSI bus 110 run compatible versions of 10 software.

As with the nodes 120, a node 220 typically includes one or more of the following: a CPU (not shown), a memory (not shown), a user interface (not shown), a co-processor (not shown), a port (not shown), a communications interface 221 and an internal bus 127. (The composition 15 of a node 220 may be altered according to the node's application in the arts of embedded systems and distributed computing.)

The memory typically includes RAM (not shown), as well as non-volatile memory (also not shown). Further, the memory typically contains layered software: Application software 223 communicates with 20 the operating system 12222, and the operating system 12222 includes an I/O subsystem 12223. The I/O subsystem 12223 communicates with the user interface, the co-processor, the ports and the communications interfaces 221 by means of the communications bus 127.

A communications interface 221, in this embodiment, may be a 25 RAID controller 221.

The communications bus 127 communicatively interconnects the CPU, memory, user interface, co-processor, port and communications interfaces 221.

To each SCSI bus 110 is connected at least one device 2061, 30 each device 2061 having a SCSI ID unique for the channel 110 to which it connects. Local disks 150 connect to respective nodes 220 as necessary.

Figure 9 illustrates the logical view of the cluster system

200 of Figure 2. The RAID controllers 221 organize the devices 2061 to appear to the hosts 220 as SCSI disks 130 on SCSI channels 110. The RAID controllers 221 thus appear to the host 120 as HBAs 221.

The computer system 200 is a shared-nothing system. That is to say, while more than one node 220 in the system 200 may have access to a device, one and only one node 220 owns and manages that device at any given time.

Figure 3 illustrates a controller 221 for a node 220 of a cluster system 200 according to the invention. A controller 221 includes a CPU 340, random access memory (RAM) 310, non-volatile memory (NVM, typically NVRAM) 320, read-only memory (ROM) 330 and I/O controllers 350 and 360. Each of the controller components 310 through 360 communicates using the internal bus 370. The I/O controller 360 communicatively connects the controller 221 to the node internal bus 127, while the I/O controller 350 communicatively connects the controller 221 to the SCSI bus 110.

In various embodiments, the RAM 310 includes a reservation table 390 and possibly a monitor process 380 (both described below).

Reservation of a logical device 130 is done by the cooperating controllers 221. On the operating system 12222's requesting the reservation of a logical device 130, the driver for the logical device 130 issues a command to the (software of) the controller 221. The controller 221 handles the reservation as per SCSI specifications.

In handling the reservation at the controller level, a controller 221 communicates reservation information to any other controllers 221, communicating via the back-end SCSI bus(es) 110. On successful communications, the reserving controller 221 completes the command to the operating system 12222.

The system 200 implements the Release() command at the level of the controllers 221 as well. The controller 221 receiving the Release() command communicates to any other controllers 221 before completing the command to the operating system 12222.

## Data Structures

A first data structure is a list of channels connecting one particular node controller 221 to other node controllers 221. The 5 instant implementation of the list is as a bit map. (The pseudo-code used herein tends to follow the grammar of the C and C++ languages.):

```
int Connected_Channels;
```

10 A second data structure is a node 220's SCSI ID, Node\_ID. Such a data structure is familiar to a routine practitioner of the SCSI art. As the described embodiment is a SCSI-2 implementation, a Node\_ID can be represented as an 8-bit character:

```
15 char Node_ID;
```

A third data structure is a list of SCSI IDs for the nodes 220. The list is here implemented as an array Node\_IDs:

```
20 char Node_IDs[];
```

A next data structure, Node, encompasses a node 220's SCSI ID and its state:

```
25 struct Node {    char Node_ID;
                    char Node_State;
                    char Node_State_Info;
                    struct Node *Next_Node;
                };
```

30

The values of the Node\_State field include: MASTER, NEW\_SLAVE, DIAGNOSED\_SLAVE, RESTARTED\_SLAVE, FAILED\_SLAVE, NEXUS\_SLAVE,

ACTIVE\_SLAVE and KILLED\_SLAVE.

The MASTER state indicates that the controller 221 of the node 220 is the master of the cluster system 200.

5 The NEW\_SLAVE state indicates that the master node controller 221 has recently recognized a node controller 221 other than itself.

The DIAGNOSED\_SLAVE state indicates that diagnostics that the master node controller 221 performed on a slave node controller 221 completed successfully.

10 The FAILED\_SLAVE state indicates that the diagnostics performed on the slave node controller 221 did not complete successfully.

The RESTARTED\_SLAVE state indicates that the master node controller 221 has directed a slave node controller 221 to reconfigure itself and restart.

15 On successful completion of negotiations, the state of a slave node controller 221 is NEXUS\_SLAVE.

The state ACTIVE\_SLAVE indicates that a slave node controller 221 has started up successfully and is awaiting negotiations with the master node controller 221.

20 The KILLED\_SLAVE state indicates that a slave node controller 221 timed out on negotiations, could not update its configuration or has been expressly killed by a node controller 221.

25 The Node\_State\_Info field gives additional information about the state of a node controller 221. The values of the Node\_State\_Info field include: Starting\_Up, Waiting\_For\_Negotiations, Negotiating, Failed\_Diagnostics, Timed\_Out\_On\_Negotiations, Changing\_Configuration, Failed\_Communications\_Test, Completed\_Startup and Rebuilding.

“Starting\_Up” indicates that a node controller 221 is starting up.

30 “Waiting\_For\_Negotiations” indicates that a slave node controller 221 is awaiting a Node\_Negotiate() command from the master node controller 221.

“Negotiating” indicates that a slave node controller 221 is just about to start or is now negotiating its configuration with the master node controller 221.

5 “Failed\_Diagnostics” indicates that a slave node controller 221 failed a diagnostic.

“Timed\_Out\_On\_Negotiations” indicates that a slave node controller 221 timed out waiting for the master node controller 221 to send it a *Node\_Negotiate()* command.

10 “Changing\_Configuration” indicates that a slave node controller 221 is changing its configuration and restarting.

“Failed\_Communications\_Test” indicates that a slave node controller 221 has failed a test of communications between the master node controller 221 and the slave node controller 221.

15 “Completed\_Startup” indicates that a node controller 221 has completed its startup process.

“Rebuilding” indicates that a node controller 221 is aware of a rebuild in progress in the cluster system 200.

20 A next data structure is a list of nodes 220 and their respective states. In one embodiment, this list is a linked list, *Nodes*, of nodes 220:

```
struct Node *Nodes
```

25 A next data structure is a reservation table 390. A reservation table 390 cross-references logical devices (logical disks 130, in the cluster system 900) with a reserving SCSI device, if any (hosts 220, in the cluster system 900). Table 1 illustrates one embodiment of a reservation table 390.

30 In Table 1, the node controller 221 with SCSI ID 7 has reserved the logical disk 130 with ID 4. The node controller 221 with SCSI ID 6 has reserved the logical disk 130 with ID 0.

A next data structure is the SCSI-specified 10-byte vendor-

specific command with opcode 20h. Table 2 describes the SCSI format of the 20h command. As with any SCSI command, the first byte (byte 0) is an opcode and the last byte (here, byte 9) is a control byte. The remaining fields of a 20h-opcode command are vendor specific.

5 An embodiment of this invention uses the 10-byte 20h-opcode vendor-specific command to effect communications between nodes 220. More particularly, the invention uses the 20h-opcode vendor-specific command to create commands issued, recognized and executed by cooperating controllers 221 to:

- 10 • check channel connections between two nodes 220 in a cluster system 200 (Node\_Cable\_Check());
- read data from another node controller 221 in a cluster system 200 (Node\_Read());
- write data to another node controller 221 in a cluster system 200 (Node\_Write());
- 15 • update configuration information in another node controller 221 in a cluster system 200 (Node\_Negotiate());
- force a node controller 221 into a suspension mode (Node\_Kill());
- update the reservation tables 390 of other nodes 220 in a cluster system 200 while reserving a logical device (Node\_Reserve());
- 20 • remove a logical-device reservation from other nodes 220 in a cluster system 200 (Node\_Release());
- inform another controller 221 in a cluster system 200 of the receipt of a command to rebuild a logical device 130 (Node\_Inform\_Rebuild\_Received());
- 25 • inform a master node controller 221 in a cluster system 200 of the specifics of a command to rebuild a logical device 130 (Node\_Inform\_Master\_Rebuild\_Received());
- instruct a node controller 221 in a cluster system 200 to start a rebuild of a logical device 130 (Node\_Rebuild());
- 30 • inform a master node controller 221 in a cluster system 200 that a slave node controller 221 has responsively started a rebuild of

5 a logical device 130 (Node\_Inform\_Rebuild\_Started());

- periodically inform a node controller 221 in a cluster system 200 of the progress of a rebuild (Node\_Inform\_Status());
- inform a node controller 221 in a cluster system 200 of the completion of a rebuild of a logical device 130 (Node\_Inform\_Rebuild\_Completed()); and
- inform a slave node controller 221 that a rebuild of all logical devices 130 associated with a new drive is complete (Node\_Inform\_Rebuilt\_All()).

10

Each of these commands is described in turn below.

Table 3 describes the format of the Node\_Cable\_Check() command. Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the logical unit number (LUN), and the remaining bits of byte 1 are unused (here, reserved).  
15 Byte 2 specifies that the command is Node\_Cable\_Check(). The remaining bytes are unused.

Table 4 describes the format of the Node\_Read() command. Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Read(). Bytes 3-6 specify the starting address, with bit 0 of byte 3 as the LSB and bit 7 of byte 6 as the MSB. Bytes 7 and 8 are respectively the LSB and MSB of the number of bytes to read.  
25

Table 5 describes the format of the Node\_Write() command. Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Write(). Bytes 3-6 specify the starting address, with bit 0 of byte 3 as the LSB and bit 7 of byte 6 as the MSB. Bytes 7 and 8 are respectively the LSB and MSB of the number of bytes to read.  
30

Table 6 describes the format of the Node\_Negotiate() command.

Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Negotiate(). Bytes 3-8 are unused.

5 Table 7 describes the format of the Node\_Kill() command.

Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Kill(). Bytes 3-8 are unused.

10 Table 8 describes the format of the Node\_Reserve() command.

Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Reserve(). Bytes 3-8 are unused.

15 Table 9 describes the format of the Node\_Release() command.

Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Release(). Bytes 3-8 are unused.

20 Table 10 describes the format of the

Node\_Inform\_Rebuild\_Received() command. Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Inform\_Rebuild\_Received(). Bytes 3-8 are

25 unused.

Table 11 describes the format of the  
Node\_Inform\_Master\_Rebuild\_Received() command. Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Inform\_Master\_Rebuild\_Received(). Bytes 3-8 are unused.

Table 12 describes the format of the Node\_Rebuild() command.

Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Rebuild(). Bytes 3-8 are unused.

5 Table 13 describes the format of the Node\_Inform\_Rebuild\_Started() command. Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Inform\_Rebuild\_Started(). Bytes 3-8 are unused.

10 Table 14 describes the format of the Node\_Inform\_Status() command. Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Inform\_Status(). Bytes 3-8 are unused.

15 Table 15 describes the format of the Node\_Inform\_Rebuild\_Completed() command. Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that the command is Node\_Inform\_Rebuild\_Completed(). Bytes 3-8 are 20 unused.

Table 16 describes the format of the Node\_Inform\_Rebuilt\_All() command. Byte 0 is the opcode, and byte 9 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are unused. Byte 2 specifies that 25 the command is Node\_Inform\_Rebuilt\_All(). Bytes 3-8 are unused.

Table 17 describes the format of a Heartbeat() command, in one embodiment the SCSI six-byte Test Unit Ready() command. Byte 0 is the opcode, and byte 5 is the control byte. The three most significant bits of byte 1 specify the LUN, and the remaining bits of byte 1 are 30 unused. Bytes 2-4 are unused.

A next data structure is the SCSI Inquiry() command with opcode 12h, described in Table 17. Byte 0 is the opcode, and byte 5 is

the control byte. The three most significant bits of byte 1 specify the LUN, the least significant bit is the Enable Vital Product Data bit, and the second-least significant bit is the Command Support Data bit. The remaining bits of byte 1 are unused. Byte 2 specifies the page 5 code/opcode. Byte 3 is unused. Byte 4 specifies the data length.

A next data structure is the standard SCSI Inquiry() data format, described in Table 19. As the standard Inquiry() data contains numerous fields, only a subset of which is of interest here, only that subset is described herein. The three most significant bits of byte 0 10 are the peripheral qualifier, and the remaining bits of byte 0 are the peripheral device type. The two most significant bits of byte 2 and the next three most significant bits of byte 2 indicate whether the device supports the ISO IS-9316 or the ECMA-111 version of the SCSI standard. The four least significant bits of byte 3 indicate the format of the data 15 following. Byte 4 specifies how many additional bytes of information follow. Bytes 8-15 indicate the manufacturer's name in ASCII. Bytes 16-31 specify the product's name in ASCII. Bytes 32-35 specify the product's version number in ASCII. Bytes 36-55 are vendor specific.

A next data structure is a mailbox as follows:

```
20          struct mailbox{
21              struct SYS_DRIVES logical_drives;
22
23              /* MAX_No_CHANNELS is the maximum number of channels
24                  possibly attached to a node.
25                  MAX_No_TARGETS is the maximum number of targets
26                  possible on a given channel. */
27              struct P_D_E physical_drives [MAX_No_CHANNELS]
28                                [MAX_No_TARGETS];
29
30          struct CONFIG configuration;
```

};

5 The logical\_drives member of the mailbox data structure reflects the state of any logical drives in the system 200. Likewise, the physical\_drives data structure reflects the state of the physical drives in the system 200.

These data structures and their use are explained further below.

## Methods, Processes and Protocols

## -- Firmware-Implemented SCSI Commands

In supporting the logical devices (130, in the cluster system 15 900) the controllers 221 implement the following commands in firmware: Reserve(), Release(), Reset\_Device(), Test\_Unit\_Ready(), Inquiry(), Read\_Capacity(), Read\_FUA\_DPO() and Write\_FUA\_DPO(). All of the above commands are well known in the SCSI art, save for the Read\_ and Write\_FUA\_DPO() commands. (“FUA” is “force unit access.” “DPO” is 20 “disable page out.” Both terms are well known in the SCSI art.) These commands may be implemented in the 60h-7Fh SCSI reserved command group.

Reservation of a logical device 130 is done by the controllers 221. On the operating system 12222's requesting the reservation of a logical device 130, the driver for the logical device 130 issues a DCMD command to the (software of) the controller 221. The controller 221 handles the reservation as per SCSI specifications.

In handling the reservation at the controller 221 level, a controller 221a communicates reservation information to any other controllers 221. The controllers communicate via the back-end SCSI bus(es) 110. On successful communications, the reserving controller 221 completes the command to the operating system 12222.

The system 200 implements the Release() command at the level

of the controllers 221 as well. The controller 221 receiving the Release() command communicates to any other controllers 221 before completing the command to the operating system 12222.

The new Read\_ and Write\_FUA\_DPO() commands are the SCSI ten-  
5 byte Read() and Write() commands with the DPO and FUA bits. In one embodiment, the support for the DPO involves doing nothing: The controllers 221 support only write throughs. Reads with the FUA bit set are treated as cache misses.

10           -- Node-to-Node Communications

The nodes 220 in the cluster system 200 communicate via the back-end SCSI bus(es) 110. A node controller 221 uses for node-to-node communications any channel 110 that the Connected\_Channels bit map notes as connected and available. The channel 110 that a node controller 221  
15 is currently using is herein termed the "node-communications channel."

Figures 11A-11C compose a flowchart illustrating flow of control in a node controller 221 processing SCSI bus communications: When a node controller 221 receives the Node\_Cable\_Check() command, step 1160, the node controller 221 returns the channel number on which it  
20 received the command, step 1165.

When a node controller 221 receives the Node\_Read() command, step 1150, the node controller 221 returns the specified count of bytes, starting at the physical address specified, step 1155.

When a node controller 221 receives the Node\_Write() command,  
25 step 1140, the node controller 221 writes the specified count of bytes, starting at the specified physical address, step 1145.

A Node\_Write() command may send a completion status only after the data is written to the media. The Node\_Read() and Node\_Write()  
30 commands support the disable-page-output (DPO) and force-unit-access (FUA) bits of the SCSI specification.

Figure 13 is a flowchart illustrating the flow of control in a node 220 processing a Node\_Negotiate() command. When a node controller

221 receives the Node\_Negotiate() command, step 1170, the node controller 221 scans the bus 110, step 1305, to determine whether the specified configuration is acceptable, step 1310. If not, the node controller 221 returns a configuration-mismatch error, step 1315. If an update to 5 EEPROM fails, steps 1320 and 1325, the node controller 221 returns an EEPROM-write-failure error, step 1315. If an update to NVM fails, steps 1330 and 1335, the node controller 221 returns a NVM-write-failure error, step 1315.

If no configuration mismatch occurs, step 1310, if the writes 10 to EEPROM and NVM succeed, steps 1320 through 1335, and if the accepted configuration differs from that the node controller 221 had earlier, step 1340, the node controller 221 informs the master node controller 221, step 1345, and restarts itself, step 1350.

In one embodiment, the mailbox data structure is the format 15 for the data exchanged.

When a node controller 221 receives the Node\_Kill() command, step 1130, the node controller 221 goes into a suspension mode, step 1135. (In suspension mode, a controller 221 awaits human intervention to correct an anomalous situation.)

20 **Figure 10** is a flowchart illustrating the flow of control in a node processing a SCSI Reserve() command. **Figure 12** is a flowchart illustrating the flow of control in a node 220 processing a Node\_Reserve() command.

When a node controller 221 receives the Node\_Reserve() 25 command, step 1120, the node controller 221 checks its reservation table 390 to see whether the specified logical device has already been reserved, step 1205. If no entry is found for the specified logical device, the node controller 221 makes the reservation for the node 220 sending the command, step 1210, and returns with GOOD status, step 1215.

30 If, however, an entry exists, step 1205, the command terminates, returning a reservation-conflict error, step 1220.

When a controller 221 receives a SCSI Reserve() command, it

checks its reservation table 390 to determine whether some node controller 221 has already reserved the logical device specified in the command, step 1005. If the specified logical device is not already reserved, the controller 221 make a reservation entry, step 1015. The 5 controller 221 then issues a Node\_Reserve() command(s) to the other node(s) 220 in the cluster system 200 to reserve the specified logical device, step 1040. In so doing, the controller 221 determines whether any other node controller 221 has reserved the specified logical device (in the other node controller 221's reservation table 390), step 1020. 10 If no node controller 221 has previously reserved the specified logical device, step 1205, the other node controller 221 reserves the specified logical device on behalf of the controller 221, step 1210, and the controller 221 successfully completes the Reserve() command, step 1030.

If another node controller 221 has reserved the specified 15 logical device, step 1020, and if the node controller 221 processing the Reserve() command has higher priority (that is to say in one embodiment, if the ID of the node controller 221 processing the Reserve() command is higher than the ID of the node controller 221 reserving the logical device), step 1050, the node controller 221 waits an interval of time, 20 step 1045, and checks the table 390 again, step 1040. If an entry is still found, step 1020, the controller 221 returns from the Reserve() command with a reservation-conflict error, step 1035.

A situation may occur where two node controllers (say, 221a and 221b) receive Reserve() commands at the same time for the same 25 logical device. The node controller 221a makes an entry in its reservation table 390a, step 1015, and attempts to reserve on the node controller 221b, steps 1040 and 1020. Substantially simultaneously, the node controller 221b makes an entry in its reservation table 390b, step 1015, and attempts to reserve on the node controller 221a, steps 1040 and 30 1020. Now each node controller 221 finds an entry in the other node controller 221's table 390. Here both node controllers 221a and 221b could effect a less than optimal solution: both returning a reservation-

conflict error.

A more desirable way to resolve this situation is described below, involving higher- and lower-priority nodes 220. In one embodiment, priority is determined by relative SCSI IDs. Given any two node controllers 221 on the same SCSI bus 110, one has a higher SCSI ID than the other. Say node controller 221a has the higher SCSI ID. It is then the higher-priority node controller 221. Whenever the higher-priority node controller 221a finds an entry in the other lower-priority node controller 221b's reservation table 390b, step 1020, and recognizes that it (node controller 221a) has higher priority, step 1050, it waits for a predetermined amount of time, step 1045, and checks the lower-priority node controller 221b's table 390b again, step 1040.

If the higher-priority node controller 221a still finds an entry, step 1020, it responds with a reservation-conflict error to the Reserve() command, step 1035. However, if it does not find the entry on the later try; step 1020, it makes the reservation on the node controller 221b, step 1040, and successfully returns from the Reserve() command, step 1030.

Meanwhile, the node controller 221b finds a conflict, step 1020. Being lower priority, step 1050, it removes the entry from its own reservation table 390b, step 1035. When the higher-priority node controller 221a checks the table 390b again, step 1040, it does not find a reservation for the specified logical device, step 1020, and makes the reservation on the subordinate node controller 221b, step 1040.

Figure 19 is a flowchart illustrating flow of control in a node controller 221 processing a SCSI Release() command. On a Release() command, a controller 221 removes the entry from its table 390 and from the table(s) 390 of the other node controller(s) 221, step 1905. The controller 221 returns GOOD status, step 1910. (The controller 221 returns a GOOD status even if the command is for a logical device that is not reserved on the controller 221.)

When a node controller 221 receives a Node\_Release() command,

step 1110, the node controller 221 removes any reservation for the specified logical device from its reservation table 390, step 1105.

When a node controller 221 receives a Heartbeat() command, step 1190, the node controller 221 returns a GOOD status, step 1196, if 5 its status is in fact good, step 1193. Otherwise, the node controller 221 returns a CHECK CONDITION, step 1195.

When a slave node controller 221 receives a Node\_Inform\_Rebuild\_Received() command, step 11A0, the node controller 221 sets its Node\_State\_Info to Rebuilding, step 11A5.

When a master node controller 221 receives a Node\_Inform\_Master\_Rebuild\_Received() command, step 11B0, the master node controller 221 checks whether it has reserved the specific logical device 130, step 1605. (Figure 16 is a flowchart illustrating flow of control in a master node controller 221 on receipt of a 15 Node\_Inform\_Master\_Rebuild\_Received() command.) If so, the master node controller 221 starts the rebuild of the specific logical device 130, step 1610, and informs the other node controller(s) 221 that it successfully started the rebuild, step 1615. The master node controller 221 sends subsequent, periodic progress updates to the other 20 controller(s) 221, step 1620. On completion of the rebuild, the master node controller 221 sends a notification of that completion as necessary, step 1625. The master node controller 221 checks whether any more logical devices 130 need to be rebuilt, step 1630. If so, it starts the rebuild of the next logical device 130, step 1610. On the completion of 25 the rebuild of all logical drives 170, the master node controller 221 sends a notification of that meta-completion as necessary, step 1640. The master node controller 221 changes its Node\_State\_Info back to its previous value in preparation for accepting further rebuild commands, step 1645.

30 Figure 18 is a flowchart illustrating flow of control in a node controller processing a Node\_Rebuild() command. When a slave node controller 221 receives a Node\_Rebuild() command, step 11C0, the slave

node controller 221 starts the rebuild of the specific logical device 130, step 1805, and informs the other node controller(s) 221 that it successfully started the rebuild, step 1810. The slave node controller 221 sends subsequent, periodic progress updates to any other 5 controller(s) 221, step 1815. On completion of the rebuild, the slave node controller 221 sends a notification of that completion, step 1820.

When a node controller 221 receives a `Node_Rebuild_Start()` command, step 11D0, the node controller 221 notes the specifics of the rebuild, including its current status, step 11D5. In the event of, say, 10 the failure of the node performing the rebuild or a query from a user regarding the status of the rebuild, the node 221 has sufficient state information to take over or provide status of the rebuild.

Similarly, when a node controller 221 receives a `Node_Inform_Status()` or `Node_Inform_Rebuild_Completed()` command, step 15 11E0 or 11F0, the node controller 221 updates its status for the rebuild, step 11E5.

When a slave node controller 221 receives a `Node_Inform_Rebuilt_All()` command, step 11G0, the slave node controller 221 updates its status for the rebuild, now changing its `Node_Info_Status` 20 back to its previous value, ready to accept a next rebuild command from a node 220, step 11H0. (The master node controller 221 does not receive a `Node_Inform_Rebuilt_All()` as only it issues that particular command.)

#### -- Node Startup

25 **Figures 4 and 5A-5C** illustrate the startup protocols, according to one embodiment of the invention. **Figure 4** illustrates the startup protocol for a master node controller 221, and **Figures 5A-5C** illustrate the startup protocol for a slave node controller 221.

When a SCSI device controller 221 starts up, it identifies 30 and creates a bit map for all channels 110 connected to itself, step 4A5. The controller 221 also starts all of its devices 130, step 405, and identifies any other nodes 220 in the cluster system 200, step 410, in

one embodiment by scanning the back-end SCSI channel(s) 110. (It may find no other node controller 221.) The controller 221 enters the SCSI IDs of all identified nodes 220 (including itself) in the list of SCSI IDs, Node\_IDs, step 415.

5 In a single-node cluster system 200, the controller 221 does not identify any other node controller 221. The controller 221 sets itself to be the master of the cluster system 200, steps 425, 430 and 435. (Step 440 is essentially a NOOP.) The controller 221 processes **COD information**, step 445, updating its configuration as necessary, and  
10 completes its startup, step 4A0. It sets its Node\_State\_Info to Completed\_Startup.

When multiple nodes 220 exist ab initio in the cluster system 200, each of the nodes 220 identifies all of the other nodes 220, step 410. All of the nodes 220 use the same predetermined algorithm to  
15 determine the master/slave relationship between themselves, step 425.

In one embodiment, each node controller 221 reads the state of all the other nodes 220. As all nodes 220 are just starting up, no node controller 221 is yet the master node controller 221. The nodes 220 compare SCSI IDs, and all nodes 220 set the master node controller 221  
20 to be the node controller 221 with the highest SCSI ID, steps 425. All other nodes 220 are slaves.

The master node controller 221 initially sets the Node\_State for itself as Master and its Node\_State\_Info as Starting\_Up, step 435. The master node controller 221 also sets the Node\_State of all other  
25 nodes 220 as NEW\_SLAVE and the Node\_State\_Info of each other node controller 221 to Waiting\_For\_Negotiations, step 440.

The state MASTER applies to a node controller 221 designated as the master. Further, the master node controller 221 maintains the NEW\_SLAVE, DIAGNOSED\_SLAVE, RESTARTED\_SLAVE, FAILED\_SLAVE and NEXUS\_SLAVE  
30 states regarding slave nodes 220.

A slave node controller 221 maintains the ACTIVE\_SLAVE and KILLED\_SLAVE states regarding itself.

Accordingly, on a slave node controller 221, the node controller 221 sets the Node\_State for itself as ACTIVE\_SLAVE with Node\_State\_Info of Waiting\_For\_Negotiations, step 505. A slave node controller 221 also sets the Node\_State for the master node controller 221 to MASTER with the corresponding Node\_State\_Info set to Starting\_Up, step 510.

10 A node 220 joining the cluster system 200 starts up with a power on or reset. These events cause the controller of the node 220 to issue a SCSI bus reset that in turn causes the master node controller 221 to scan the back-end SCSI channels again. The master node controller 221 at this point detects the presence of the new node controller 221, comparable to step 410. The controller 221 of the master node 220 enters the new node 220's SCSI ID in the node identifier list Node\_ID, comparable to step 415.

15 The new node controller 221 also detects the master node controller 221 at the same time. Each of the nodes 220 reads the state of the other 220. The new node controller 221 thus learns that the master node controller 221 is in fact the master of the cluster system 220. The new node controller 221 accordingly makes itself a slave.

20 The controller 221 then creates a list, Nodes, listing all the nodes 220 identified in the Node\_IDs list, step 420. Each entry in the Nodes list contains the node ID and the state of the node controller 221.

25 The master node controller 221 then processes COD information, step 445, while the slave node controller 221 skips processing COD information and waits for negotiations from the master node controller 221, steps 520, 525 and 530.

30 At initialization, the node controller 221 sets the node-communications channel 110 used for node-to-node communications, step 4B0.

-- The Monitor Process

On completion of its startup, the master node controller 221 periodically runs a monitor process 380. (The periodicity has been 5 determined empirically to be about every 5 seconds.) Figures 8A-8C compose a flowchart illustrating actions of the monitor process 380.

The monitor process 380 searches the node identifier list Node\_ID to see if a new node 220 entered the cluster system 200 since its last check, step 805. If the monitor process 380 finds a new node 10 220, step 810, it creates an entry in the nodes list Nodes, step 815. The new entry has the Node\_ID of the newly found node 220, a Node\_State of NEW\_SLAVE and a Node\_State\_Info of Waiting\_For\_Negotiations.

The master node controller 221 performs diagnostics on the new slave node, step 860. These may include cable-check and memory 15 tests. If the slave node performs the diagnostics successfully, step 865, the master node controller 221 sets the Node\_State of the slave node controller 221 in the master node controller 221 to DIAGNOSED\_SLAVE with the corresponding Node\_State\_Info set to Negotiating, step 870.

If diagnostics fail, step 865, the master node controller 221 20 changes the Node\_State of the slave node controller 221 to FAILED\_SLAVE with the corresponding Node\_State\_Info set to Failed\_Diagnostics, step 875. The monitor process 380 completes its tasks, delays (i.e., sleeps on an alarm, dies and resurrects on a timer, loops, etc.), step 8A0, and begins again to search for new nodes 220, step 805.

25 On the master node controller 221 determining the slave node controller 221 to be a FAILED\_SLAVE, that slave node controller 221 times out waiting for negotiations with the master node controller 221, steps 520, 525 and 530. On the timeout, the slave node controller 221 changes its Node\_State to KILLED\_SLAVE with Node\_State\_Info of 30 Timed\_Out\_On\_Negotiations, step 535. Finally, the slave node controller 221 puts itself into a suspension mode, step 540.

On successful diagnostics, step 865, the master node

controller 221 updates the configuration and other information of the slave, step 8B0. The master node controller 221, using the Negotiate() command, requests the slave controller 221 to verify the information, step 880. The slave node controller 221 verifies the configuration, step 5 545, scanning the bus again. If the new configuration is acceptable, step 550, the slave node controller 221 updates its EEPROM, NVM and other configuration structures as necessary, step 555.

The Negotiate() command may fail either because the master node controller 221's configuration does not match that of the slave node 10 or because the write to the configuration structures failed. In either of these failing situations, step 885, the master node controller 221 puts the slave node controller 221 in the FAILED\_SLAVE state with Node\_State\_Info giving the reason for failure, step 890. Further, the slave node controller 221 changes its Node\_State to KILLED\_SLAVE with 15 Node\_State\_Info giving the reason, steps 550, 560 and 570, and goes into a suspension mode, step 540.

On successful negotiation, steps 550 and 560, if the new configuration of the slave node controller 221 is different from what the slave node controller 221 had earlier, step 575, the slave node 20 controller 221 so informs the master node controller 221, step 580, and restarts itself, step 585. The master node controller 221 changes its Node\_State of the slave node controller 221 to RESTARTED\_SLAVE with Node\_State\_Info set to Changing\_Configuration, step 895. The master node controller 221 then resumes negotiations after a specified interval of 25 time, steps 8E0 and 860 through and 880.

If the new configuration of the slave node controller 221 is the same as what the slave node controller 221 had earlier, step 575, the slave node controller 221 completes its startup. The master node controller 221 changes its Node\_State of the slave node controller 221 to 30 NEXUS\_SLAVE and the corresponding Node\_State\_Info to None, steps 8C0 and 8C5. The master node controller 221 completes its tasks. Both controllers 221 are ready to accept commands.

(Of course, where the master node controller 221 detects more than one new node 220, the master node controller 221 may perform the above steps in parallel for two or more of the new nodes 220.)

Figure 17 is a flowchart illustrating flow of control in a 5 node controller 221 monitoring another controller 221. The master node controller 221 monitors the slave node controllers 221 and vice versa, issuing Heartbeat() commands, step 1710. (For efficiency reasons, slave node controllers 221 do not monitor each other.) As long as a 10 Heartbeat() command is successful, step 1715, the master node controller and slave node controller 221 are in nexus. (In one embodiment, the Heartbeat() command fails four consecutive times before a node controller 221 recognizes (i.e., formally acknowledges) a Heartbeat() failure and 15 proceeds to change another node controller 221's status to FAILED\_SLAVE, steps 1705, 1720 and 1725.)

15 If a Heartbeat() command fails, step 1715, the failure-detecting node controller 221 tries to determine the cause of failure. The Heartbeat() can fail for either of the following two reasons: (1) the other node 220 reset or powered off or (2) the cable of the node-communications channel 110 failed. If cable failure caused the 20 Heartbeat() failure, the node controller 221 kills itself, but if the other node 220 is cycling, the node controller 221 attempts communications on the next connected channel.

More particularly, the failure-detecting node controller 221 determines the first drive 2061 online on the node-communications channel 25 110, step 1745. If no drive 2061 is online on the channel, step 1730, the node sets the node-communications channel to the next connected channel, as indicated by the Connected\_Channels bit map, and performs the Heartbeat() command on the next connected channel, step 1755.

Where a drive 2061 is available and online on the node- 30 communications channel 110, step 1730, the node controller 221 sends a Test Unit Ready() command to that drive 2061, step 1745. Where the Test Unit Ready() command is successful, step 1735, the cable presumptively

has not failed. Accordingly, the node controller 221 detecting the Heartbeat() failure changes the Node\_State of the failing node controller 221 to FAILED\_SLAVE. The failure-detecting node controller 221 then attempts to kill the failing node controller 221 with a Node\_Kill() 5 command, step 1740.

Where a drive 2061 is available and online on the node-communications channel 110, step 1730, but the Test Unit Ready() command is unsuccessful, step 1735, the cable presumptively has failed. Accordingly, the node controller 221 detecting the Heartbeat() failure 10 kills itself, step 1750.

If the failing, killed node controller 221 is the master node controller 221, then the surviving node controllers 221 arbitrate for the role of the master (as in step 428, for example). If the former master controller 221 resets and attempts to re-join the cluster system 200 15 again, it assumes the role of a slave controller 221 in the protocols and processes described herein. As a corollary, the former-slave, now-master controller 221 treats the former-master controller 221 as a slave node controller 221.

Because the detection of the failure of a node 220 can take a 20 relatively long time in terms of computer systems (in one embodiment, fifteen seconds), a node 220 can fail and recover before another node 220 has time to detect the failure. If the failure is not detected and correctly dealt with, the recovering node 220 may not be able to establish a nexus with the computer system 200.

25 To avoid this failure-to-detect-failure scenario, a controller 221 of a node 220 of a computer system 200 detects other nodes 220 in the system 200 after a bus reset by scanning the bus(es) 110 and obtaining SCSI IDs as described herein. In addition, the node controller 221 sends an Inquiry() command to a detected node 220. From the 30 Inquiry() data returned, the scanning node controller 221 can determine whether the detected node controller 221 is in the process of booting up (i.e., recently failed) or is already up (i.e., already in nexus). In

one embodiment, with the Inquiry() data fields set as indicated in Table 20, byte 36 (vendor-specific data) is a flag indicating whether the node 220 is in the booting process.

A node 220 that failed to detect the failure of another node 5 220 (due, for example, to the quickness of the power cycle or the reset) can parse the Inquiry() data to determine whether any node 220 is booting up and thus not in nexus. The inquiring node 220 compares its Node\_State for the other node 220 with the Inquiry() data and thus determines whether the other node 220 had failed. If the other node had failed, the 10 inquiring node 220 sets its Node\_State and Node\_State\_Info for the other node as NEW\_SLAVE, Waiting\_For\_Negotiations. Where the recovering node 220 was the master node 220, the new master node 220 sets its own state as MASTER and proceeds to establish nexus with the recovery node 220.

In addition to detecting new nodes 220, the monitor process 15 380 also checks existing connections, step 825 (for example, directing the master node controller 221 to write to a slave node controller 221's memory 310). If the check times out, steps 830 through 840, the master node controller 221 changes its Node\_State of the slave node controller 221 to FAILED\_SLAVE and the corresponding Node\_State\_Info to 20 Failed\_Communications\_Test, step 845. The master node controller 221 then puts the slave node controller 221 in a suspension mode, step 850. This step 850 may or may not succeed, depending on the nature of the communications failure.

If a node controller, say, 221a, determines that another node 25 controller, say, 221b, is not functioning satisfactorily, then the node controller 221a kills the other node controller 221b, using the Node\_Kill() command. If the killed node controller 221b is the master node controller 221, then it loses that status. The surviving node controllers (including controller 221a) arbitrate to become the master 30 node controller 221 (as in step 425, for example).

Figure 6 is a state diagram giving the states transitions a master node controller 221 makes regarding a slave node controller 221.

Figure 7 is a state diagram giving the state transitions a slave node controller 221 makes regarding itself.

-- Configuration Changes

5 Every time the configuration changes in the master node controller 221, the master node controller 221 updates itself, the SCSI devices 2061 and the slave node(s) 220 with the changed configuration. If the configuration changes on a slave node controller 221, the slave node controller 221 updates itself and informs the master node controller 10 221. The master node controller 221 then updates itself, any other slave node controllers 221 and all the devices 2061. In one embodiment, only the master node controller 221 writes configurations to the devices 2061.

A controller 221 detecting cable failure can potentially kill all the drives on the failing channel and distribute the resulting 15 configuration by negotiation, causing I/O from other node controllers 221 to these killed drives to also then fail. Before, however, another controller 221 enters negotiations to update the configuration of the system 200 in response to a Negotiate() command, it attempts to determine whether the configuration change is due to cable failure. If a cable 20 failure caused the configuration change, it refrains from negotiating a new configuration for the system 200 and instead kills itself.

More particularly, on receiving a Negotiate() command, the node controller 221 sets the channel for testing to the first connected channel, step 1510. The controller 221 then checks that channel for 25 testing, step 1515. If the channel passes the test, step 1520, the controller checks for more channels, step 1525. Where there are more channels, the controller 221 sets up the next connected channel as the current channel for testing, step 1530, and repeats the check of the channel for testing, step 1515.

30 Where the channel does not pass the test, step 1520, the controller determines whether any physical drives are online on the channel failing the channel check, step 1540. Where there are no online

drives on the channel, the controller 221 proceeds to set up the next connected channel, if any (step 1525), as the current channel for testing, step 1530, and repeats the check of the channel for testing, step 1515.

5           Where there is a drive online on the channel failing the channel check, step 1540, the controller gets the first drive online on the channel, step 1545, and performs a Test Unit Ready() command on the drive, step 1550. If the Test Unit Ready() command succeeds, step 1555, the controller proceeds with the negotiation of the configuration change, 10 step 1560, secure that a cable failure did not cause the configuration change.

Where, however, the Test Unit Ready() command fails, the node controller 221 kills itself, step 1565, having determined that the configuration change is due to a cable failure.

15           Where there are no more connected channels, step 1525, the controller 221 proceeds to negotiate the change in configuration of the system 200, step 1570, having determined that the change is not due to a cable failure.

20           -- Rebuild of a Logical Device

Figure 14 is a flowchart illustrating flow of control in a node controller 221 receiving a rebuild command. A node controller 221 receiving a rebuild command from its node 220 checks whether any rebuild is in process in the cluster system 200, step 1405. If no rebuild is in 25 progress, the node controller 221 sends Node\_Inform\_Rebuild\_Received() commands to any other node controllers 221, step 1410. If the node \_controller 221 is a slave controller 221, step 1415, then the node controller 221 additionally sends a Node\_Inform\_Master\_Rebuild\_Received() command, step 1420. Independent of its master/slave status, the node 30 controller sets its Node\_State\_Info to Rebuilding, step 1425.

On receipt of a Node\_Inform\_Rebuild\_Received() command, step 11A0, a node controller 221 sets its Node\_State\_Info to Rebuilding, step

11A5. In one embodiment, the cluster system 200 permits only one rebuild at a time. Any node controller 221 in a Rebuilding state rejects any subsequent rebuild command from its node 220, steps 1405 and 1430.

Figure 16 is a flowchart illustrating flow of control in a master node controller 221 on receipt of a Node\_Inform\_Master\_Rebuild\_Received() command. When the master node controller 221 receives the specifics of a rebuild command (directly from its node 220 or indirectly from a slave node controller Node\_Inform\_Master\_Rebuild\_Received() command), step 11B0, the master node controller 221 checks whether it owns the specific logical device 130, step 1605. If so, the master node controller 221 starts the rebuild of the specific logical device 130, step 1610. It sends out Node\_Inform\_Rebuild\_Started() commands, step 1615, and subsequent, periodic Node\_Inform\_Status() commands to any other controllers 221, step 1620. On completion of the rebuild, the master node controller 221 sends Node\_Inform\_Rebuild\_Completed() commands, step 1625, and checks whether any more logical devices 130 need to be rebuilt, step 1630. If so, it starts the rebuild of the next logical device 130, step 1605. On the completion of the rebuild of all logical drives dependent on a failed physical drive 2061, step 1630, the master node controller 221 sends Node\_Inform\_Rebuilt\_All() commands as necessary, step 1640. The master node controller 221 changes its Node\_State\_Info back to its previous value in preparation for accepting further rebuild commands, step 1645.

Where the master node controller 221 does not have the reservation for the specific logical device, step 1605, it determines which node controller 221 does have the reservation. It sends a Node\_Rebuild() command to that determined node controller 221, step 1650.

On receipt of a Node\_Rebuild() command, step 11C0, a slave node controller 221 starts the rebuild of the specific logical device, step 1805. It also sends Node\_Inform\_Rebuild\_Started() commands and subsequent, periodic Node\_Inform\_Status() commands to the other node controller(s) 221, steps 1810 and 1815. On completion of the rebuild,

the slave node controller 221 sends Node\_Inform\_Rebuild\_Completed() commands as necessary, step 1820.

The preceding description also applies to a node controller 221's detecting and initiating a rebuild without the intercession of a 5 node 220, for example, on the automatic detection of the removal and insertion of a physical drive 2061 constituting a logical drive 130.

While the description above is in terms of SCSI-2 and a cluster system of two nodes, a routine practitioner in the art will readily realize the invention's applications in a later SCSI environment. 10 for example, SCSI-3 or in a cluster system having 3 or more nodes. Indeed, the invention now being fully described, many changes and modifications that can be made thereto without departing from the spirit or scope of the appended claims will be apparent to one of ordinary skill in the art. The application of the invention in an external-RAID- 15 controller environment such as that of Figure 20 and/or with a reservation table 390 as in Table 21 are examples.

This specification incorporates by reference all publications and patent applications mentioned herein, to the same extent if the specification had specifically and individually incorporated by reference 20 each such individual publication or patent application.

**Table 1:** Reservation Table

TARGET SCSI ID	RESERVING INITIATOR SCSI ID
0	6
1	(NONE)
...	...
4	7

**Table 2:** 20h-Opcde Vendor-Specific SCSI Command

Bit Byte	7	6	5	4	3	2	1	0
0	Opcode = 20h							
1	Vendor-Specific							
2	Vendor-Specific							
3	Vendor-Specific							
4	Vendor-Specific							
5	Vendor-Specific							
6	Vendor-Specific							
7	Vendor-Specific							
8	Vendor-Specific							
9	Control Byte							

**Table 3: Node\_Cable\_Check()**

Bit Byte	7	6	5	4	3	2	1	0					
0	Opcode = 20h												
1	Logical Unit Number			Reserved									
2	DCMD Opcode - Node_Cable_Check												
3	Reserved												
4	Reserved												
5	Reserved												
6	Reserved												
7	Reserved												
8	Reserved												
9	Control Byte												

Table 4: Node\_Read()

Bit Byte	7	6	5	4	3	2	1	0						
0	Opcode = 20h													
1	Logical Unit Number		Reserved											
2	DCMD Opcode - Node_Read													
3														
4	Physical Address													
5														
6	MSB													
7	Count (LSB)													
8	Count (MSB)													
9	Control Byte													

**Table 5: Node\_Write()**

Bit Byte	7	6	5	4	3	2	1	0					
0	Opcode = 20h												
1	Logical Unit Number			Reserved									
2	DCMD Opcode - Node_Write												
3	LSB												
4	Physical Address												
5													
6	MSB												
7	Count (LSB)												
8	Count (MSB)												
9	Control Byte												

**Table 6: Node\_Negotiate()**

Bit Byte	7	6	5	4	3	2	1	0					
0	Opcode = 20h												
1	Logical Unit Number			Reserved									
2	DCMD Opcode - Node_Negotiate												
3	Reserved												
4	Reserved												
5	Reserved												
6	Reserved												
7	Reserved												
8	Reserved												
9	Control Byte												

**Table 7:** Node\_Kill()

**Table 8: Node\_Reserve()**

Bit Byte	7	6	5	4	3	2	1	0				
0	Opcode = 20h											
1	Logical Unit Number				Reserved							
2	DCMD Opcode - Node_Reserve											
3	Reserved											
4	Reserved											
5	Reserved											
6	Reserved											
7	Reserved											
8	Reserved											
9	Control Byte											

**Table 9: Node\_Release()**

Bit Byte	7	6	5	4	3	2	1	0				
0	Opcode = 20h											
1	Logical Unit Number				Reserved							
2	DCMD Opcode - Node_Release											
3	Reserved											
4	Reserved											
5	Reserved											
6	Reserved											
7	Reserved											
8	Reserved											
9	Control Byte											

**Table 10: Node\_Inform\_Rebuild\_Received()**

Bit Byte	7	6	5	4	3	2	1	0					
0	Opcode = 20h												
1	Logical Unit Number			Reserved									
2	DCMD Opcode - Node_Inform_Rebuild_Received												
3	Reserved												
4	Reserved												
5	Reserved												
6	Reserved												
7	Reserved												
8	Reserved												
9	Control Byte												

**Table 11: Node\_Inform\_Master\_Rebuild\_Received()**

Bit Byte	7	6	5	4	3	2	1	0						
0	Opcode = 20h													
1	Logical Unit Number		Reserved											
2	DCMD Opcode - Node_Inform_Master_Rebuild_Received													
3	Reserved													
4	Reserved													
5	Reserved													
6	Reserved													
7	Reserved													
8	Reserved													
9	Control Byte													

**Table 12: Node\_Rebuild()**

Bit Byte	7	6	5	4	3	2	1	0					
0	Opcode = 20h												
1	Logical Unit Number			Reserved									
2	DCMD Opcode - Node_Rebuild												
3	Reserved												
4	Reserved												
5	Reserved												
6	Reserved												
7	Reserved												
8	Reserved												
9	Control Byte												

**Table 13: Node\_Inform\_Rebuild\_Started()**

Bit Byte	7	6	5	4	3	2	1	0				
0	Opcode = 20h											
1	Logical Unit Number				Reserved							
2	DCMD Opcode - Node_Inform_Rebuild_Started											
3	Reserved											
4	Reserved											
5	Reserved											
6	Reserved											
7	Reserved											
8	Reserved											
9	Control Byte											

**Table 14: Node\_Inform\_Status()**

Bit Byte	7	6	5	4	3	2	1	0						
0	Opcode = 20h													
1	Logical Unit Number		Reserved											
2	DCMD Opcode - Node_Inform_Status													
3	Reserved													
4	Reserved													
5	Reserved													
6	Reserved													
7	Reserved													
8	Reserved													
9	Control Byte													

**Table 15: Node\_Inform\_Rebuild\_Completed()**

Bit Byte	7	6	5	4	3	2	1	0						
0	Opcode = 20h													
1	Logical Unit Number		Reserved											
2	DCMD Opcode - Node_Inform_Rebuild_Completed													
3	Reserved													
4	Reserved													
5	Reserved													
6	Reserved													
7	Reserved													
8	Reserved													
9	Control Byte													

**Table 16: Node\_Inform\_Rebuilt\_All()**

Bit Byte	7	6	5	4	3	2	1	0					
0	Opcode = 20h												
1	Logical Unit Number			Reserved									
2	DCMD Opcode - Node_Inform_Rebuilt_All												
3	Reserved												
4	Reserved												
5	Reserved												
6	Reserved												
7	Reserved												
8	Reserved												
9	Control Byte												

**Table 17: Heartbeat()**

Bit Byte	7	6	5	4	3	2	1	0					
0	Opcode = 00h												
1	Logical Unit Number			Reserved									
2	Reserved												
3	Reserved												
4	Reserved												
5	Control Byte												

**Table 18: Inquiry()**

Bit Byte	7	6	5	4	3	2	1	0			
0	Opcode = 12h										
1	Logical Unit Number			Reserved			CmdDt	EVDP			
2	Page Code										
3	Reserved										
4	Data Length										
5	Control Byte										

Table 19: Standard SCSI Inquiry() Data

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier							
1	RMB							
2	ISO							
3	AEN							
4	TIO (rmTsk)							
5	Reserved							
6	Reserved (Port)							
7	Reserved (DualP)							
8-15	Reserved (MChngr)							
16-31	Reserved (ARQ)							
32-35	Sync							
36-55	Link							
56-95	Res.							
95-	(TmDis)							
	Manufacturer							
	Product							
	Revision							
	Vendor Unique							
	Reserved							
	Vendor Unique							

**Table 20:** Vendor-Specific Inquiry() Data

Field	Value
Peripheral Qualifier	3
Device Class	0
ISO	02
Data Format	02
Additional Length	1Fh
Manufacturer	MYLEX
Product	DAC960
Version	F/W Version

**Table 21:** Alternate Reservation Table

<b>TARGET SCSI ID</b>	<b>LOGICAL UNIT ID</b>	<b>RESERVING INITIATOR SCSI ID</b>
0	0	6
	1	(NONE)
	...	...
	7	7
1	0	(NONE)
	1	(NONE)
	...	...
	7	7
...	...	...
4	0	(NONE)
	1	(NONE)
	...	...
	7	(NONE)

WHAT IS CLAIMED IS:

1           1. A method for managing access to an I/O device, said  
2 method comprising:  
3           communicatively coupling first and second nodes, having  
4           respective first and second bus controllers, and a logical I/O  
5           device by means of a bus and said first and second controllers;  
6           receiving on said first controller a request to reserve said  
7           logical I/O device; and  
8           communicating by means of said bus from said first to said  
9           second controller a reservation request for said logical I/O device  
10           for execution by said second controller, in response to said  
11           receiving.

1           2. The method of claim 1, further comprising the step of  
2           reserving said logical I/O device for said first node within  
3           said second controller, in response to said reservation request  
4           communication.

1           3. The method of claim 2, wherein said step of reserving  
2           comprises:  
3           determining whether said logical I/O device is already  
4           reserved within said second controller;  
5           communicating a response, indicating failure to reserve said  
6           logical I/O device, to said first node when said logical I/O device  
7           is already reserved; and  
8           otherwise, reserving said logical I/O device for said first  
9           node within said second controller, and communicating to said first  
10           node a response indicating success in reserving said logical I/O  
11           device.

1                   4. The method of claim 3, further comprising the steps of:  
2                   receiving said response to said communicated reservation  
3                   request;

4                   aborting the method for managing access when said response  
5                   indicates failure to reserve and said first controller is  
6                   subordinate to said second controller; and

7                   otherwise, delaying and communicating again a reservation  
8                   request for said logical I/O device when said response indicates  
9                   failure to reserve and said first controller is dominant to said  
10                  second controller;

11                  otherwise, responding, indicating success, to said received  
12                  reservation request.

1                   5. The method of claim 1, wherein said step of  
2                   communicatively coupling comprises  
3                   communicatively coupling said first and second nodes and said  
4                   logical I/O device depending from a multi-logical-device, third  
5                   controller by means of said bus and said first and second  
6                   controllers.

1                   6. The method of claim 1, wherein after said step of  
2                   receiving and before said step of communicating, the following steps are  
3                   performed:

4                   in response to said reservation request, determining whether  
5                   said logical I/O device is already reserved within said first  
6                   controller, and aborting said method for managing access when said  
7                   logical I/O device is already reserved; and

8                   otherwise, reserving said logical I/O device for said first  
9                   node within said first controller.

1                   7. A computer-readable medium for data storage wherein is  
7.

1 located a computer program for causing a first node in a computer system,  
2 having a first bus controller, to manage access to a logical I/O device  
3 in said computer system by

4 receiving on said first controller a request to reserve said  
5 logical I/O device; and

6 communicating by means of a bus from said first controller to  
7 a second controller of a second node a reservation request for said  
8 logical I/O device for execution by said second controller, in  
9 response to said receiving.

1 8. The computer-readable medium of claim 7, wherein said  
2 computer program causes access management by further

3 reserving said logical I/O device for said first node within  
4 said second controller, in response to said reservation request  
5 communication.

1 9. The computer-readable medium of claim 8, wherein said  
2 step of reserving in said computer program comprises

3 determining whether said logical I/O device is already  
4 reserved within said second controller;

5 communicating a response, indicating failure to reserve said  
6 logical I/O device, to said first node when said logical I/O device  
7 is already reserved; and

8 otherwise, reserving said logical I/O device for said first  
9 node within said second controller, and otherwise, reserving said  
10 logical I/O device for said first node within said second  
11 controller, and communicating to said first node a response  
12 indicating success in reserving said logical I/O device.

1 10. The computer-readable medium of claim 7, wherein after  
2 said step of receiving and before said step of communicating in said

1                   11. A computer system comprising:  
2                   first and second nodes having respective first and second bus  
3                   controllers, said first controller comprising  
4                   the computer-readable medium of claim 7; and  
5                   a CPU, coupled to said medium, for executing said  
6                   computer program in said medium;  
7                   an I/O device; and  
8                   a bus communicatively coupling said first and second nodes  
9                   and said logical I/O device by means of said first and second  
10                  controllers.

1                    12. A method for managing access to a logical I/O device,  
2 said method comprising:  
3                    communicatively coupling first and second nodes having  
4                    respective first and second bus controllers, and a logical I/O  
5                    device by means of a bus and said first and second controllers;  
6                    receiving on said first controller a request to release said  
7                    logical I/O device; and  
8                    communicating by means of said bus from said first to said  
9                    second controller a release request for said logical I/O device for  
10                   execution by said second controller, in response to said receiving.

13. The method of claim 12, wherein before said step of

1 receiving, the following steps are performed:

2 receiving on said first controller a request to reserve said  
3 logical I/O device; and

4 communicating by means of said bus from said first to said  
5 second controller a reservation request for said logical I/O device  
6 for execution by said second controller, in response to said  
7 receiving a reservation request.

1 14. The method of claim 12, further comprising the step of  
2 releasing said logical I/O device within said second  
3 controller, in response to said release request communication.

1 15. The method of claim 12, wherein said step of

2 communicatively coupling comprises

3 communicatively coupling said first and second nodes and a  
4 logical device depending from a multi-logical-device, third  
5 controller by means of said bus and said first and second  
6 controllers.

1 16. A computer-readable medium for data storage wherein is  
2 located a computer program for causing a first node in a computer system,  
3 having a first bus controller, to manage access to a logical I/O device  
4 in said computer system by

5 receiving on said first controller a request to release said  
6 logical I/O device; and

7 communicating by means of a bus from said first controller to  
8 a second controller of a second node a release request for said  
9 logical I/O device for execution by said second controller, in  
10 response to said receiving.

1 17. The computer-readable medium of claim 16, wherein said

1 computer program manages access by further  
2 releasing said logical I/O device within said second  
3 controller, in response to said release request communication.

18. A computer system comprising:  
first and second nodes having respective first and second bus  
controllers, said first controller comprising  
the computer-readable medium of claim 16; and  
a CPU, coupled to said medium, for executing said  
computer program in said medium;  
an I/O device; and  
a bus communicatively coupling said first and second nodes  
and said logical I/O device by means of said first and second  
controllers.

19. An apparatus for managing access to a logical I/O  
device, said apparatus comprising:  
means for communicatively coupling first and second nodes,  
having respective first and second bus controllers, and a logical  
I/O device;  
means for receiving on said first controller a request to  
reserve said logical I/O device; and  
means for communicating from said first to said second  
controller a reservation request for said logical I/O device for  
execution by said second controller, in response to said receiving.

20. An apparatus for managing access to a logical I/O  
device, said apparatus comprising:  
means for communicatively coupling first and second nodes,  
having respective first and second bus controllers, and a logical  
I/O device;

1           means for receiving on said first controller a request to  
2 release said logical I/O device; and

3           means for communicating by means of said bus from said first  
4 to said second controller a release request for said logical I/O  
5 device for execution by said second controller, in response to said  
6 receiving.

1


EE EE EE EE EE EE EE EE  
EE EE EE EE EE EE EE EE

# APPARATUS AND METHOD FOR IMPLEMENTING RAID DEVICES IN A CLUSTER COMPUTER SYSTEM


## ABSTRACT OF THE DISCLOSURE


Apparatus and methods for efficiently implementing logical-device reservations in a cluster computer system. The apparatus includes 9 cooperating controllers programmed in firmware around a distributed reservation table. The apparatus manages access to a logical device, with first and second nodes with respective bus controllers communicatively coupled to each other and to a logical device by means of 13 a bus. The first controller receives a request to reserve the logical device and, in response, communicates a reservation request for the logical device over the bus to the second controller for execution by the second controller. In response to the communicated reservation request, 17 the second controller reserves the logical device for the first node. The logical device may be a logical unit depending from a multi-logical-unit controller. The apparatus further manages access by the first controller's receiving a request to release the logical device and 21 responsively communicating a release request over the bus to the second controller for execution by the second controller. The second controller responsively releases the logical device within itself.

*Figure 1 (Prior Art)*




**Figure 2**






*Figure 3*





*Figure 4*

*Figure 5A*




**Figure 5B**







*Figure 5C*



**Figure 6**


*Figure 7*






**Figure**  
**8A**


**Figure 8B**




*Figure 8C*




**Figure 9**




*Figure 10*




*Figure 11A*




*Figure 11B*




*Figure 11C*







*Figure 12*

*Figure 13*







*Figure 14*




*Figure 15*

**Figure 16**






*Figure 17*




*Figure 18*




*Figure 19*


*Figure 20 (Prior Art)*



*Figure 21 (Prior Art)*



*Figure 22 (Prior Art)*



*Figure 23 (Prior Art)*

